In review, Cognition, Technology, and Work

The Decision Ladder as an Automation Planning Tool

M. Cummings
Aeronautics & Astronautics Department
Massachusetts Institute of Technology
77 Massachusetts Ave
Cambridge, MA 02139

Stephanie Guerlain
Department of Systems & Information Engineering
University of Virginia
Charlottesville, VA 22904-4747

Communicating Author: M. Cummings

Abstract

The decision ladder was developed to model information processing activities and resultant knowledge states for control tasks. This paper shows how the decision ladder can also provide guidance on the incorporation of automation in decision support system design. Using the decision ladder as an automation planning tool, potential areas for automation were identified for the United States Navy's new Tactical Tomahawk missile supervisory control user interface. An equally important result was recognizing where automation should not be used. Testing of a low fidelity decision aid was conducted as a result of the insight gained through using the decision ladder for automation analysis, which revealed the potential for automation bias. This experiment demonstrated that higher levels of automation do not necessarily translate into improved controller performance. In the case of the Tactical Tomahawk design, the automation level for retargeting decisions was minimized.

Keywords

Decision ladder, automation, decision support system design, automation bias

1.0 Introduction

The Tomahawk missile is the U.S. Navy's premier land attack missile, and indeed, the U.S. military has declared, "Because of its long range, lethality, and extreme accuracy, *Tomahawk*® has become the weapon of choice for the U.S. Department of Defense (U.S. Navy 2000)." An earlier version of this missile was used during the Gulf War in 1991, against Al Qaida in Afghanistan, and most recently in Operation Iraqi Freedom. It is strategically invaluable since it can be fired approximately 1000 miles from its intended target with pinpoint accuracy. The U.S. Navy is in the process of designing a new version of the missile, called the Tactical

Tomahawk, which will provide battlefield commanders with the ability to redirect Tomahawk missiles in-flight. The Navy foresees that within a single "strike" (a predetermined time span in which a group of missiles and other weapons are launched for a common mission objective), a total of 128 missiles could launched. In addition, the Navy would like to implement an entirely new mission for the Tomahawk as an overhead loitering missile that circles until redirected to an emergent target. An example of an emergent target would be the emergence of surface-to-air enemy missile launch platforms whose positions are unknown until they actually begin electronic transmissions.

The implementation of the Tactical Tomahawk means that not only will battlefield commanders have more flexibility and options; it also means that a layer of human control will be needed where none previously existed. Introducing the ability for an operator to control a very fast-moving tactical weapon in the close-in combat arena as well as managing high value assets through constant replanning requires substantial cognitive contribution. To comprehensively model the decision making process required for in-flight retargeting of missiles and identity both information requirements and necessary states of knowledge, the decision ladder model approach was used. Because the Tactical Tomahawk system is revolutionary, it is not clear how in-flight retargeting of missiles decision making process can best be supported and what levels of automation will be needed to support this task. As a result of mapping the in-flight retargeting decision making process using the decision ladder, it was discovered that while not a primary objective of the decision ladder, the decision ladder was an effective automation planning aid. This paper will detail how the decision ladder was applied to the Tactical Tomahawk case, and how it was used as an automation planning tool. In addition, automation possibilities identified through this analysis were experimentally tested and the results are discussed.

2.0 The Decision Ladder

To illustrate how control tasks and actions relate to the decision-making sequences of a system operator, Rasmussen (1976) developed the decision ladder model, which represents information processing activities and subsequent states of knowledge that result when the activities are performed in a decision making process. The decision ladder framework has been applied to numerous domains to include hospitalization diagnostic sessions (Rasmussen, Pejtersen et al. 1994), thermal-hydraulic process control (Vicente 1999), production scheduling (Sanderson 1991; Higgins 2001), military domains such as broad command and control networks (Chin, Sanderson et al. 1999) and threat management for naval frigates (Chalmers, Easter et al. 2000). The decision ladder maps rather than models the structure of a decision-making process, and in doing so, helps to identify the requirements of decision-making control tasks. The decision ladder model usually focuses on expert performance to demonstrate not how a system *should* act, but a system *could* act if given a flexible problem space. The decision ladder should represent what information processes need to occur, independent of who will perform a task or how a particular control task will be accomplished (Rasmussen, Pejtersen et al. 1994). In the case of systems with computer-based decision support tools, the decision ladder represents the decision process and states of knowledge that must be addressed by the tool whether or not a computer or a human makes the decision.

Figure 1 illustrates the decision ladder that was developed to map the decision-making process that needs to occur in the in-flight retargeting of a Tactical Tomahawk missile for an emergent target. The left side of the decision ladder, the upward leg, represents the information

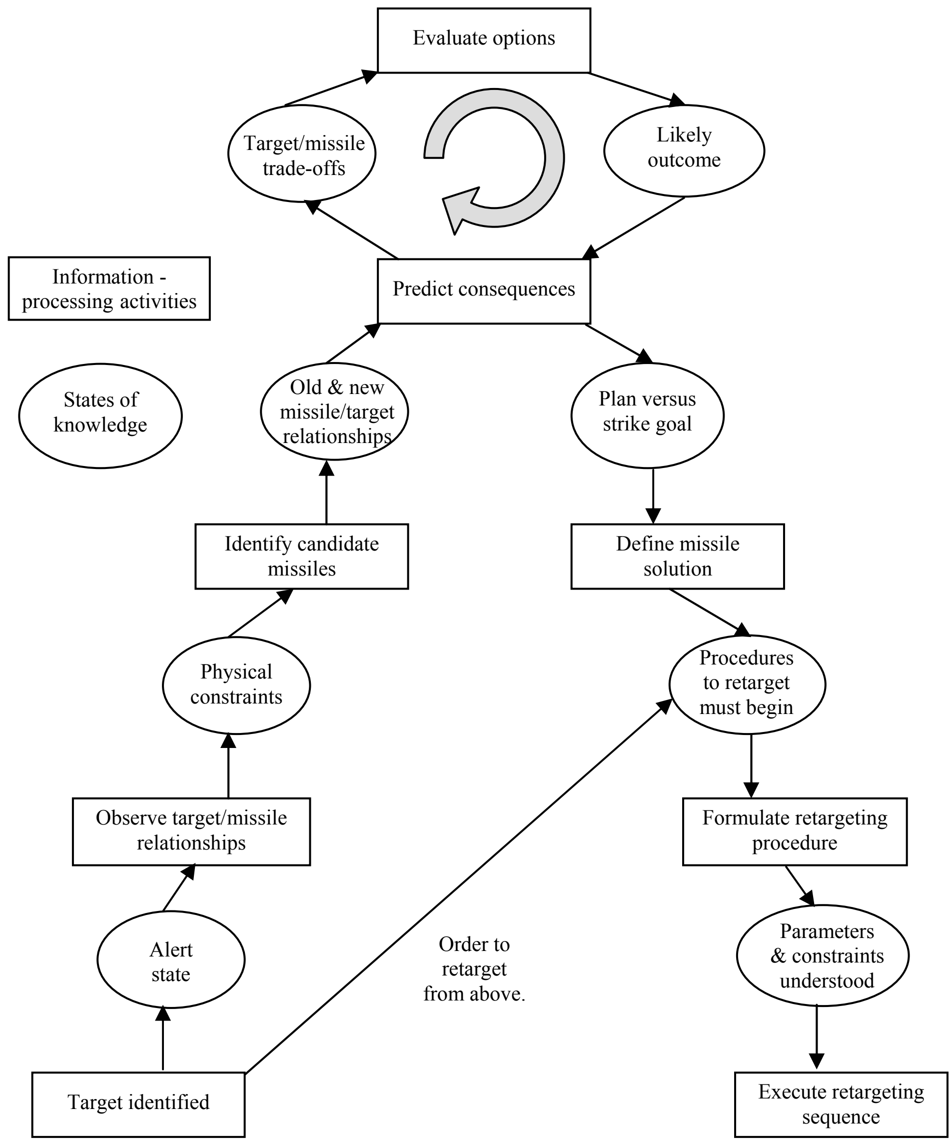

Figure 1: Initial Decision Ladder for Tactical Tomahawk Retargeting Task

processes that are required for situation analysis, the pinnacle depicts the judgment, and the downward leg symbolizes the implementation of the decision. For the Tactical Tomahawk interface missile selection/retargeting sequence, the left side of the decision ladder represents notification of an emergent target situation and the data that must be gathered to make an informed decision. The top of the decision ladder represents the optimal missile selection process, and the circular arrow depicts the iterative nature of this decision. For each retargeting scenario, the decision is affected by various factors, which are never the same due to the stochastic nature of combat. The right side of the decision ladder represents those activities and states of knowledge that are required to actually implement the decision to retarget.

While the decision ladder represents the fundamental knowledge-based decision processes, it is possible that within the ladder, an expert would be able to skip certain rungs of the ladder based on experience and familiarity with the system. Representing these "shunts, " or shortcuts makes the decision ladder a flexible tool that allows for understanding of different levels of decision-making, i.e. the decision ladder can represent both novice and expert decision-making processes (Vicente 1999) as well as reflect possible users' heuristics (Rasmussen, Pejtersen et al. 1994). When an actor takes a shortcut, various information-processing actions are bypassed but the desired results are still achieved. The decision ladder not only displays these shortcut relationships in information-processing activities, it also highlights those states of knowledge that are bypassed if a shortcut is taken. For example, in the Tactical Tomahawk interface, when a controller is notified of an emergent target, it is possible that the controller could be given orders from a higher, external authority to use a specific missile against a particular target. If this is indeed the case, then the controller no longer has to move up the ladder to evaluate all missiles weighted against various other factors. The decision of which missile to use was made by an external agency and thus the controller moves from the identification stage to the definition of missile solution phase to begin the action of retargeting. Essentially the decision-making process is removed from the controller and all that remains is to begin the actual sequence of commands to retarget.

However, human processes are not the only way shortcuts can be introduced into the decision-making process. Another manner in which various elements of the decision ladder could be bypassed is through the introduction of automation. A primary design consideration cognitive systems engineers must take into account is how automation can best support the operator and what level of automation should be introduced into a decision support system to provide human-centered automation support (Billings 1997; Parasuraman and Riley 1997; Parasuraman 2000; Parasuraman, Masalonis et al. 2000). Various levels of automation can be introduced into decision support systems, from fully automated where the operator is completely left out of the decision process to minimal levels of automation where the automation only makes recommendations and the operator has the final say. For rigid tasks that require no flexibility in decision-making and with a low probability of system failure, higher levels of automation often provides the best solution (Endsley and Kaber 1999; Kaber and Endsley in press) However, in systems like those that require decision-making in dynamic environments with many external and changing constraints, higher levels of automation are not advisable because of the inability of the automated decision aid to be perfectly reliable (Wickens 1999; Sarter and Schroeder 2001).

Although not discussed in the literature to date, we propose the decision ladder can be used as a preliminary design tool in mapping automation strategies. While the primary purpose of the decision ladder is to identify control tasks and knowledge states, because of its ability to represent the information and knowledge states bypassed through shortcuts, is can be used to aid in automation design strategies.

3.0 The Decision Ladder as an Automation Planning Tool

Mapping potential decision-making strategies is critical for designing decision support systems in conceptual development like the Tactical Tomahawk since this mapping begins to give structure and form to what was previously mere ideas and hypotheses during the work domain analysis. Since no operators or decision support tool exist for the Tactical Tomahawk system that can shed light on either novice or expert strategies, mapping possible shortcut decision paths that can occur through various levels of automation is critical for the development of a decision support system design. The decision ladder provides this initial structure through an organized framework which maps information processes and knowledge states required for decision-making tasks (Rasmussen, Pejtersen et al. 1994). Shunts in the decision ladder represent shortcuts that can be taken in a decision-making process, which bypass some information-processing actions but still achieve the desired results. These shunts can occur through both human processes as well as through the introduction of automation. Human-in-the-loop designs which employ automation for redundant, manual, and monotonous tasks while allowing operators to actively participate in the decision-making process provide not only safety benefits, but allow a human operator, and thus a system to respond more flexibly to uncertain and unexpected events (Parasuraman 2000). However, there can be measurable costs to human performance when automation is used, such as loss of situational awareness, complacency, automation bias, and skill degradation (see Parasuraman, Sheridan, and Wickens 2000 for a review.) In summary, the "black box" approach to full automation can be useful for redundant tasks that require no knowledge-based judgments, but the subsequent lack of system understanding and loss of situational awareness that full automation can cause could lead to unanticipated effects for more complex tasks.

Vicente (1999) hints at the use of the decision ladder as a guide for automation when he says the overall goal of control task analysis is to "determine how computer-based support systems could be designed to allow workers to effectively meet the challenges they face." He also suggests that analysis of potential shortcuts is critical so that they can be incorporated into the design of automation. We take the additional to step to propose that the decision ladder can be used to illustrate where automation could and should be introduced, and to what degree for the design of decision support systems. This is an especially critical step for the design of systems in conceptual development where it is clear that because of complicated technology, automation will be needed, but due to a lack of an established system or domain, degrees of automation are unclear.

Applying the levels of automation for decision and action selection originally proposed by Sheridan and Verplank (1978) and modified by Wickens et al. (1998) (Table 1) to the original decision ladder (Figure 1), a new decision ladder (Figure 2) was conceptualized for the Tactical Tomahawk retargeting task, which suggests where possible opportunities exist for automation. While other levels and scales of automation and autonomy have been proposed (Endsley and Kiris 1995; Wickens, Mavor et al. 1998; Endsley and Kaber 1999; Parasuraman, Sheridan et al. 2000), the levels in Table 1 were selected because they concentrate on decision and action, and demonstrate the more general roles and actions of both the human operator and the computer for each level of automation. These more broad categories are preferable in the conceptualization stage since the system and domain are not yet defined.

Table 1: Levels of Automation

Automation Level	Automation Description
1	The computer offers no assistance: human must take all decision and actions.
2	The computer offers a complete set of decision/action alternatives, or
3	narrows the selection down to a few, or
4	suggests one alternative, and
5	executes that suggestion if the human approves, or
6	allows the human a restricted time to veto before automatic execution, or
7	executes automatically, then necessarily informs humans, and
8	informs the human only if asked, or
9	informs the human only if it, the computer, decides to.
10	The computer decides everything and acts autonomously, ignoring the human.

The revised decision ladder demonstrates those areas of the missile retargeting decision-making process that can be automated to some degree. For example, regardless of whether a human or computer identifies the candidate missiles for an emerging target, the same information must be considered such as which missiles have enough fuel remaining to get to the target, which ones have the correct warhead etc. For a hypothetical launch of 128 missiles, it would unnecessarily cognitively overwhelm a controller to perform such a laborious search routine based on the many requirements for candidate missile selection. Hence this cognitively intense, rule-based segment of the decision-making process is an ideal candidate for at least level 7

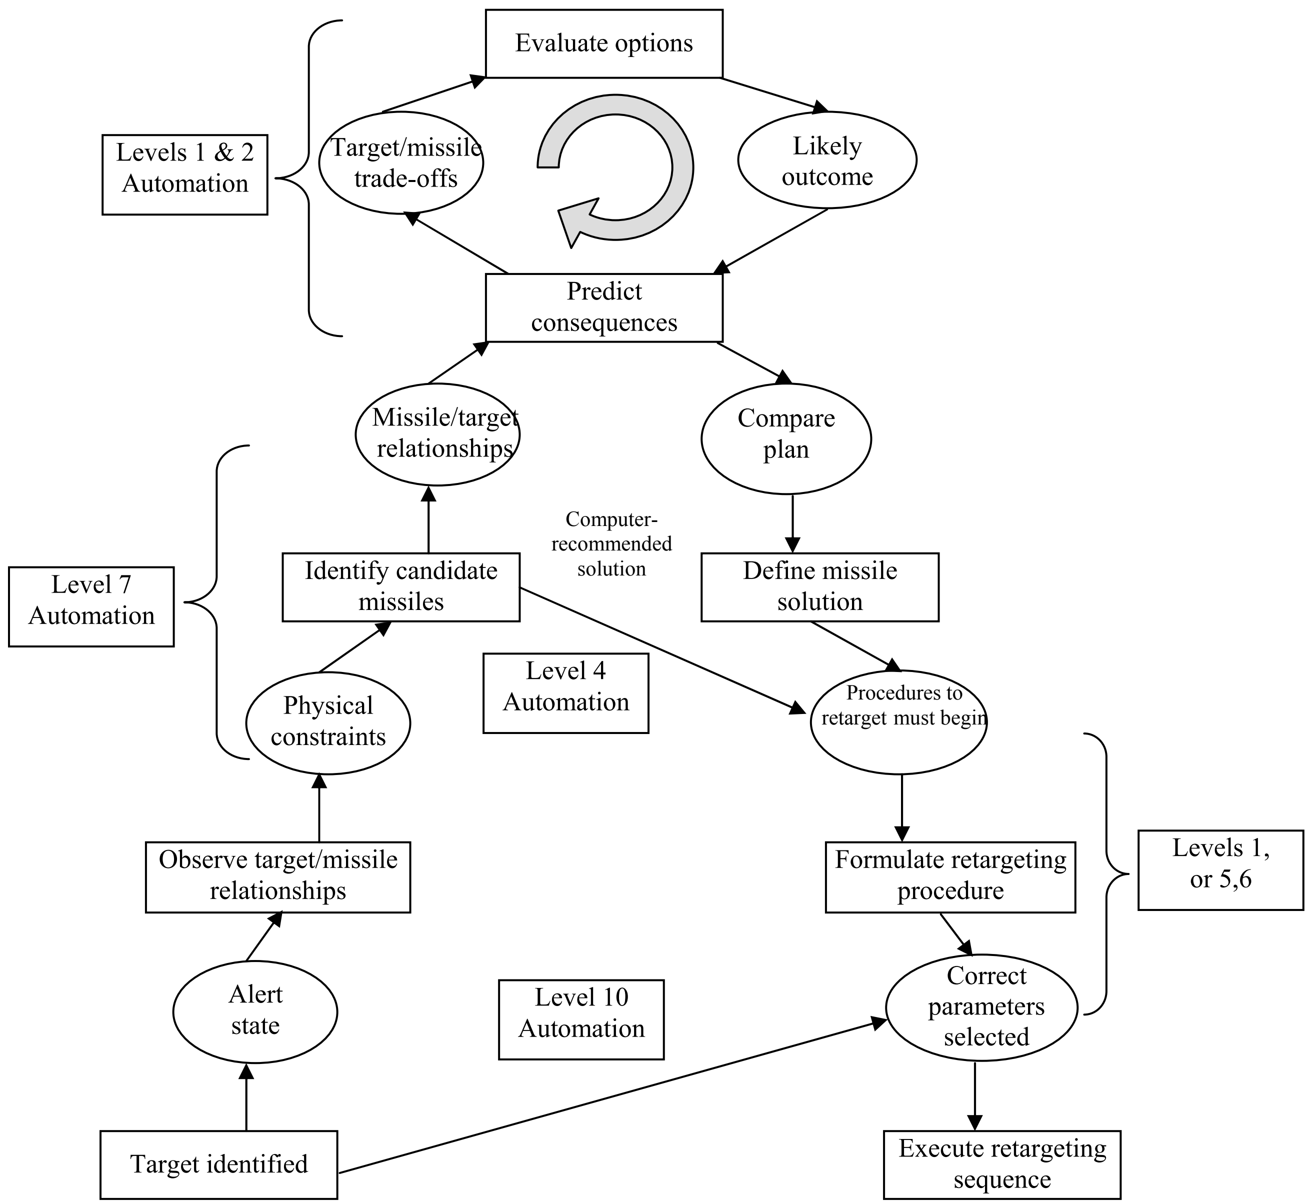

Figure 2: The Revised Tactical Tomahawk Decision Ladder

automation (automation selects the candidate missiles and informs the operator), and no knowledge states are lost that are needed in another process.

However, for the segment at the apex of the decision ladder that requires either the human or the computer to evaluate candidate missile options and predict possible outcomes, it is not clear what level of automation should be used. For the Tactical Tomahawk, the linear processes required to select candidate missiles strongly indicates the need for incorporating high levels of automation in this stage. In contrast, the iterative nature in evaluating all possible missile candidates and choosing the best missile suggest some possibilities in the incorporation of automation, but at lower levels. The primary difference between the information-processing activities of optimal missile selection as opposed to those of candidate missile selection is the ambiguous nature of an “optimal” choice.

Perhaps an even more important use of the decision ladder as an automation identification and planning tool is recognizing where automation should not be used. This critical relationship is highlighted due to the unique ability of the decision ladder to identify knowledge states and their causal and resultant activities. When specifically analyzing a decision ladder in the context of automation, a designer can more readily recognize which knowledge states would be bypassed if higher levels of automation were used. For example, if the decision to retarget missiles was fully automated, level 10 (represented by the shunt at the bottom of Figure 2), then essentially the entire decision-making process represented by the ladder is internal to the computer. Even if a lower level of automation were used that only provided recommendations for retargeting (e.g. the best 3 missiles) represented by the level 4 shunt in Figure 2, it is possible that states of knowledge are bypassed that are critical for mission success.

Higher levels of automation may also not be appropriate in situations where a computer algorithm cannot be expected to weigh all relevant facts to determine a solution. For example, when deciding which missiles are most appropriate for an emergent target, the decision maker may be informed that an intelligence agency predicts the probable appearance of another emergent target in the near future. In general, this type of projected tactical information cannot be programmed into a computer algorithm, and selecting the appropriate missile requires knowledge-based reasoning, as opposed to the rule-based activities of candidate missile selection. Because of the knowledge needed to "satisfice" in an uncertain situation, this element of the retargeting decision process would not be suited to higher levels of automation (level 6 or higher, Table 1).

When using higher levels of automation that produces shunts in the decision ladder like the level 4 or 10 shunts in Figure 2, the computer's algorithm still follows the process represented in the decision ladder. However, possible states of knowledge that the human possesses that the computer does not would not be factored into the decision. An example of this would be the intelligence case mentioned previously or a missile that is experiencing degradation in communications capabilities that could influence the decision outcome. In addition, over-automation of a system can significantly reduce situational awareness of an operator, and can degrade the effectiveness and output of the human-machine system (Endsley 1996; Parasuraman and Riley 1997; Parasuraman, Sheridan et al. 2000).

Moreover the introduction of automation can cause automation bias which occurs when operators believe that computer-generated solutions are always correct, even when presented with conflicting data that points to erroneous automation (Mosier and Skitka 1996; Parasuraman and Riley 1997).   The danger of this type of bias appearing in the Tactical Tomahawk retargeting decision is illustrated in the possible use of automation levels 5-9 for the actual retargeting process in Figure 2. The last knowledge state represented in Figure 2 is the final confirmation as to whether or not the correct missile was chosen either by the human or the computer. If level 5 or 6 were chosen for the design (see Table 1), it is possible that automation bias would be present in this knowledge state as a result of bypassing earlier knowledge states, and the controller would accept the computer's choice, regardless if it was really correct.

4.0 An Experiment to Determine an Effective Level of Automation

Prior to the development of a high fidelity prototype that would simulate the monitoring and retargeting of Tactical Tomahawk missiles, the question as to the level of automation needed for missile selection in a retargeting scenario was addressed through a low fidelity decision aid experiment. The control task analysis and the subsequent decision ladder (Figure 2) revealed that when tasked to retarget a missile, the primary decision process (the apex of the ladder) required either level 1 of automation, no computer assistance, or level 2 in which the computer provided recommendations and suggestions. To determine whether or not a computer generated recommendation would produce statistically significant results both in decision time and correctness of decision, an experiment was conducted with two versions of a low fidelity prototype, which was a printed Power Point® sketch of a decision support matrix that shows lower level data that could be inspected by a person to generate a solution in one case, and additional suggested solutions highlighted through the use of icons in another case. Testing with low fidelity paper prototypes haven been found to produce similar usability results than computer-based high-fidelity prototypes (Walker, Takayama et al. 2002).

4.1 Participants and Procedure[1]

Sixty-four University of Virginia undergraduate students enrolled in the Systems Engineering Human Machine Interface class participated as subjects. All subjects received approximately 1.5 hours of training that included a slide presentation to explain the low fidelity prototype and six practice sessions. Each subject took a qualifying quiz on symbology and matrix interpretation and only allowed to continue if they passed. The original subject N was 70 and six students were not allowed to participate due to poor quiz scores. After training, all subjects participated in two trials, with approximately a 5-minute break in between each trial. Each trial consisted of a verbally briefed scenario, a geo-spatial map of targets and missiles displayed in the front of the classroom through the use of an LCD projector, and a paper copy of a decision matrix. Figure 3 shows an example of the matrix with the additional "computer-recommended" icons that rank order the three potential solutions. Subjects were required to circle the correct answer and their decision times were recorded by fellow students with a stop watch to the nearest hundredth of a second. During one test session subjects had the icons displayed and in the other, they did not. For half of the subjects with the decision aid, intelligence information was available to the subjects that could not be quantified in the computer algorithm which made the automated recommendation the incorrect choice. Whether or not the subjects saw the decision aid test scenario first as well as the invalid recommendation session were counterbalanced across all subjects.

	ZSU-57	Weapons Depot	Truck Park	Comm Center	Bridge	Power	Sam Site	Factory	Runway
	T199U-EH	T111U-DH	T112U-DL	T113S-DM	T114U-DH	T115U-DH	T116U-DH	T117S-DL	T118U-DM
	1215 **1230**	1215 **1230** 1245	1220 **1235** 1250	1225 **1240** 1255	1230 **1245** 1300	1235 **1250** 1305	1245 **1300** 1315	1300 **1315** 1330	1315 **1330** 1345
LM033U-DL	(10:00)	(10:05)	D 19:58		(8:15)				
LM037S-DL								D 1:06:41	
RM034U-DH	(6:00)				(7:05)	D 35:13	(1:25)		(17:30)
RM036U-DM	(9:00)					(2:22)	(3:14)		D 1:10:37
FF ALLOC	0 P L	1	0	1	1	2	2	0	0
RTG/L ALLOC	0	0	1	0	0	1	0	1	1
Total Req	1	1	1	1	1	3	2	1	1
Shortage	1	0	0	0	0	0	0	0	0

Figure 3: Tactical Tomahawk Low Fidelity Decision Matrix

[1] This research protocol was approved by the UVA Social Sciences Institutional Review Board.

4.2 Apparatus

The low fidelity prototype used in this experiment is represented in Figure 3. It is a decision matrix that provides subjects with all missiles that can be redirected in a given strike as well as all targets that were preplanned for the missile strike. All missiles that can be retargeted are listed in the leftmost column and across the top are all targets in a strike. Shaded cells with a D represent a current missile/target pairing. If a cell has no numbers in it, it means the missile cannot be redirected to the target. If a number appears in parentheses, it means that missile is a potential candidate and the number is the time remaining to make the decision to retarget.

The second column in Figure 3 represents an emergent target (known to the user because the column is highlighted in red) that was assigned to the controller via electronic data link. If in a trial with automated recommendations, a circle appears next to each of the candidate solutions, the remaining decision time that graphically depicts the level of fit between the emergent target's requirements and the capabilities of the associated missile. A white circle represents a poor fit and a completely black circle indicates the best fit. The fit is determined by a predetermined computer algorithm that takes into account time to impact, the pre-assigned priorities of both the missile and target, as well as the ability of a missile to loiter and how much loiter time a missile has remaining.

4.3 Results

Two dependent variables, decision time and decision accuracy, were measured in this experiment to determine if automation level 1 or 2 significantly affected subject performance and if the use of a computer-generated recommendation promoted automation bias. The independent variables were whether or not the subject had the computer's recommendation and if the recommendation was correct or not. The decision time dependent variable was analyzed using a 2x2 repeated measures MANOVA because each subject experienced a test session with and without the decision aid. Whether or not a subject had the decision aid yielded significant results for decision time, $p < .001$, $F(1,62) = 37.1$ and a partial eta^2 of .374 ($\alpha = .05$). Those subjects with the decision aid did make decisions significantly faster than those without the decision aid. Whether or not the computer-generated recommendation was correct had no effect on decision time and there were no significant interactions.

However, speed of decision alone is not a sufficient metric for gauging performance, especially when incorrectly redirecting a million dollar missile can have severe consequences. To determine whether or not the decision aid produced measurable automation bias, accuracy of answers was analyzed. Because the independent and dependent variables are dichotomous, the McNemar test was used. The McNemar test is a nonparametric test that uses the chi-square distribution for two related dichotomous variables. For the decision aid independent variable, the McNemar test was not significant for the decision accuracy dependent variable, $p = .222$, so whether or not the subject had the decision aid did not significantly affect answer correctness. However, for the case in which subjects had the decision aid and the computer-generated recommendation was incorrect, subjects' accuracy of answers was significantly different than if the recommendation was accurate, $p = .024$. Figure 4 graphically depicts the differences in right and wrong answers for decision aiding scenarios in which the recommendation was correct and incorrect. Subjects generally selected the correct response when presented with correct computer-generated recommendations, however, when presented with a recommendation that conflicted with other instructions, the number of incorrect responses increased.

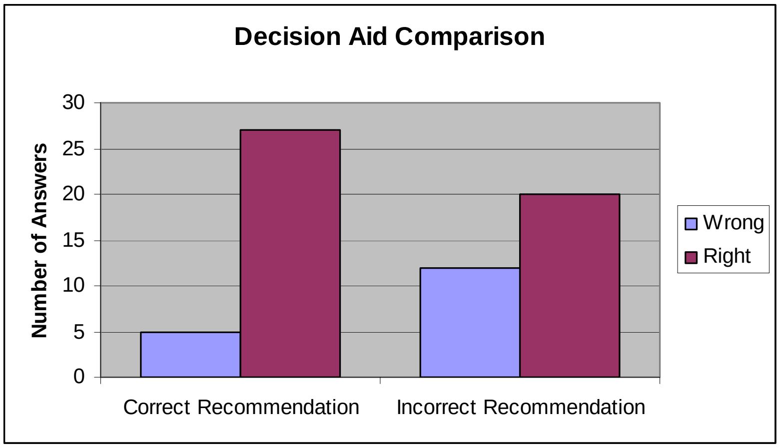

Figure 4: Decision Aiding and Answer Accuracy

4.3 Discussion

The decision aid clearly helped subjects make faster decisions but it did not necessarily help them to make correct decisions. While the McNemar test results do not establish cause and effect for answer accuracy, they suggest that the aid induces a bias and should be redesigned or not used. Other studies have demonstrated the human tendency to increasingly rely on computer-based recommendations, even though the recommended solutions are not always correct. Mosier and Skitka (1996) demonstrated that humans have a tendency to rely upon automated recommendations and pay less attention to contradictory information (see also Guerlain, Smith et al. 1996). In another study examining the effectiveness of computer-generated recommendations on pilots' decisions to counter in-flight icing problems, the results were ambiguous. When the computer provided accurate advice, pilots with the aid performed better than pilots without the aid. However, when the computer's advice was erroneous, those people without a decision aid outperformed those with one (Sarter and Schroeder 2001) Because of these results, Sarter and Schroeder recommend that unless decision aids are perfectly reliable, status displays (which merely present information) should be used instead of command displays (which recommend appropriate action.) In addition, as illustrated in a flight-planning tool, just because a human is in the loop for decision making does not mean that better solutions will be found (Layton, Smith et al. 1994), so the trade-off between the human and automation is not always clear. This inherent ambiguity is why using the decision ladder to begin to map automation strategies is useful. If designers recognize early on where automation could occur and what levels would be appropriate, potential problem human-interaction areas can be avoided later in the design.

Because of the ambiguous preliminary decision aiding results for the Tactical Tomahawk, it was decided that the high fidelity prototype would incorporate a combination of level 7 automation for finding candidate missiles (deciding which "cells" in the matrix should be filled in with a decision time), and level 1 automation for selection of the retargeting missile, i.e. the decision matrix would only show all available candidate missiles but not provide any computer-based recommendations. The Tactical Tomahawk monitoring and retargeting system is in conceptual development and no users, experts or otherwise, exist to observe or mine for knowledge. Once potential decision strategies and subject responses are gathered through testing of the high fidelity prototype, the design of a level 2 computer-generated recommendation for missiles selection will be revisited. Because of the general lack of established data and these test results that indicate the current design could promote automation bias, we decided to initially keep automation for the critical retargeting task at the lowest level possible. This more conservative design approach may be more appropriate when designing a novel decision support system (especially for a weapon control interface) until decision strategies can be developed through practice and simulations, and the risk of automation bias is more fully understood.

5.0 Conclusion

Identification of automation issues early in the design process is critical for a complex system such as the Navy's Tactical Tomahawk, which involve humans who must integrate temporal and spatial elements, as well as solve problems, manage assets, and perform contingency planning in a high workload environment. To this end, it was discovered that the decision ladder was not only useful in modeling the information processing activities and resultant states of knowledge for in-flight missile retargeting, the decision ladder also provided guidance on the incorporation of automation. While the use of the decision ladder as a tool for planning automation in a design may be purely academic for existing systems, for an extremely complex revolutionary sociotechnical system like the Tactical Tomahawk interface, the decision ladder/automation relationships provided a framework for an initial design strategy. Careful analysis of the decision ladder can reveal areas of a decision-making process that can be automated to varying degrees, but equally important, the decision ladder also highlights states of knowledge that are critical and should not be made invisible through automation.

The use of the decision ladder to evaluate areas for varying degrees of automation for the Tactical Tomahawk led to the development of a low fidelity decision aid to address an automation design problem. As the experiment with the low fidelity decision aid demonstrated, higher levels of automation do not necessarily translate into improved operator performance. Thus, perhaps an even more important result of the use of the decision ladder as a planning tool is recognizing where increasing levels of automation should not be used. When analyzing a decision ladder in the context of automation, a designer can more readily recognize which knowledge states would be lost if higher levels of automation were used. The decision ladder has the unique ability to demonstrate the impact on knowledge states of various levels of automation for systems in conceptual development with no established users or system infrastructure. Since automation in decision support systems can be a dual-edged sword, it is better to understand the issues early in the conceptual design phase than to have to correct for potentially fatal design flaws in the future.

6.0 Acknowledgments

This research was supported by the Naval Surface Warfare Center, Dahlgren Division through a grant from the Office of Naval Research, Knowledge and Superiority Future Naval Capability Program. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NSWC/DD or ONR.

References

Billings, C. E. (1997). Aviation Automation: The Search For A Human-Centered Approach. Hillsdale, N.J, Lawrence Erlbaum Associates.

Chalmers, B. A., J. R. Easter, et al. (2000). Decision-Centred Visualisations for Tactical Decision Support on a Modern Frigate. Command and Control Research and Technology Symposium, Monterey, CA.

Chin, M., P. Sanderson, et al. (1999). Cognitive Work Analysis of the Command and Control Work Domain. Command and Control Research and Technology Symposium.

Endsley, M. (1996). Automation and Situation Awareness. Automation and Human Performance. M. Mouloua. Mahwah, New Jersey, Lawrence Erlbaum Associates, Inc.**:** 163-181.

Endsley, M. and E. Kiris (1995). "The out-of-the-loop performance problem and level of control in automation." Human Factors **37**(2): 381-394.

Endsley, M. R. and D. B. Kaber (1999). "Level of automation effects on performance, situation awareness and workload in a dynamic control task." Ergonomics **42**(3): 462-492.

Guerlain, S., P. Smith, et al. (1996). "Dealing with brittleness in the design of expert systems for immunohematology." Immunohematology **12**: 101-107.

Higgins, P. G. (2001). Architecture and Interface Aspects of Scheduling Decision Support. Human Performance in Planning and Scheduling. J. Wilson, Taylor and Francis**:** 245-279.

Kaber, D. B. and M. R. Endsley (in press). "The Effects of Level of Automation and Adaptive Automation on Human Performance, Situation Awareness and Workload in a Dynamic Control Task." Theoretical Issues in Ergonomics.

Layton, C., P. J. Smith, et al. (1994). "Design of a cooperative problem-solving system for en-route flight planning: An empirical evaluation." Human Factors **36**: 94-119.

Mosier, K. L. and L. J. Skitka (1996). Human Decision Makers and Automated Decision Aids: Made for Each Other? Automation and Human Performance: Theory and Applications. M. Mouloua. Mahwah, New Jersey, Lawrence Erlbaum Associates, Inc.**:** 201-220.

Parasuraman, R. (2000). "Designing automation for human use: empirical studies and quantitative models." Ergonomics **43**(7): 931-951.

Parasuraman, R., A. J. Masalonis, et al. (2000). "Fuzzy signal detection theory: Basic postulates and formulas for analyzing human and machine performance." Human Factors **42**(4): 636-659.

Parasuraman, R. and V. Riley (1997). "Humans and Automation: Use, Misuse, Disuse, Abuse." Human Factors **39**(2): 230-253.

Parasuraman, R., T. B. Sheridan, et al. (2000). "A Model for Types and Levels of Human Interaction with Automation." IEEE Transactions on Systems, Man, and Cybernetics **30**(3): 286-297.

Rasmussen, J. (1976). Outlines of a hybrid model of the process plant operator. Monitoring behavior and supervisory control. G. Johannsen, Plenum.

Rasmussen, J., A. M. Pejtersen, et al. (1994). Cognitive Systems Engineering. New York, J. Wiley.

Sanderson, P. M. (1991). "Towards the model human scheduler." International Journal of Human Factors in Manufacturing **1**(3): 195-219.

Sarter, N. B. and B. Schroeder (2001). "Supporting decision making and action selection under time pressure and uncertainty: The case of in-flight icing." Human Factors **43**: 573-583.

U.S. Navy (2000). Tomahawk Cruise Missile, U.S. Navy. **2002**.

Vicente, K. (1999). Cognitive Work Analysis. Toward safe, productive, and healthy computer-based work. Mahwah, NJ, Erlbaum.

Walker, M., L. Takayama, et al. (2002). High-fidelity or low-fidelity, paper or computer? Choosing attributes when testing web prototypes. Humans Factors and Ergonomics Society 46th Annual Meeting, Baltimore, MD.

Wickens, C. D. (1999). Automation in Air Traffic Control: The Human Performance Issues. Automation Technology and Human Performance: Current Research and Trends. M. Mouloua. Mahwah, NJ, Lawrence Erlbaum Associates.

Wickens, C. D., A. Mavor, et al. (1998). The future of air traffic control: Human operators and automation. Washington DC, National Academy Press.